Minimally-Supervised Attribute Fusion for Data Lakes

Karamjit Singh, Garima Gupta, Gautam Shroff, and Puneet Agarwal

TCS Research, New-Delhi, India

January 5, 2017

arXiv:1701.01094v1 [cs.DB] 4 Jan 2017

Abstract

Aggregate analysis, such as comparing country-wise sales versus global market share across product categories, is often complicated by the unavailability of common join attributes, e.g., category, across diverse datasets from different geographies or retail chains, even after disparate data is technically ingested into a common *data lake*. Sometimes this is a missing data issue, while in other cases it may be inherent, e.g., the records in different geographical databases may actually describe different product 'SKUs', or follow different norms for categorization. Record linkage techniques, such as [3] can be used to automatically map products in different data sources to a common set of *global* attributes, thereby enabling federated aggregation joins to be performed. Traditional record-linkage techniques are typically unsupervised, relying textual similarity features across attributes to estimate matches. In this paper, we present an ensemble model combining minimal supervision using Bayesian network models together with unsupervised textual matching for automating such 'attribute fusion'. We present results of our approach on a large volume of real-life data from a market-research scenario and compare with a standard record matching algorithm. Finally we illustrate how attribute fusion using machine learning could be included as a data-lake management feature, especially as our approach also provides confidence values for matches, enabling human intervention, if required.

1 Introduction

Traditional business intelligence is rapidly evolving to adopt modern big-data analytics architectures based on the concept of a 'data lake', where, rather than first integrating multiple historical data from diverse sources into a common star schema via extraction-transformation-load operations, the datasets are maintained in their raw form. This leads to a number of challenges; for example, dealing with incongruous join keys between different datasets.

In this paper, we focus on a problem of fusion of information about consumer products, such as sales, market share, etc., which is spread across disparate databases belonging to different organizations, across which a product is *not* identifiable via a common key. For example, a *Global database (DB)* might track overall market-share of global product categories. On the other hand, each *Local DB* might track sales data within a geography using local-product-ids along with other characteristics, but *not* the global category-id. As a result, an analytical task such as comparing the sales of product categories within each geograpy against their global market share becomes difficult due to the lack of a natural join attribute between the databases. (Note that the same product might be characterized using different attributes in different countries, including also textual *description* of products entered manually by retailers, e.g., for carbonated drinks it usually contains information of brand, size, material used, packaging etc.)

One way to perform analysis across disparate databases is by mapping records in each *Local DB* to their corresponding global attributes (e.g., 'category' in the example above). However, preparing such mappings is a huge manual and complicated task because: a) The cardinality (number of possible values) of local and global characteristics varies from tens to thousands, and b) Uncertainty in the semantics of local characteristics of the same product from different geographies, leading to confusion in identifying the product category, even by human annotators.

Our aim is to help reduce cost of the operational process of creating and maintaining such global references by reducing manual workload via automation via modern data-lake architecture that include automated fusion of federated databases. Our goal is to either make high confidence predictions, or abstain from making any prediction so that such records can be sent to human annotators. We want to minimize the number of such abstentions while maximizing the precision of the predictions.

Attribute Fusion using Record Matching: Consider two databases (see Figure 1): a) *Local DB(L)* with each product l having local characteristics $L_1, L_2, ..., L_M$, e.g., flavor, brand, etc., and retailer descriptions (D_i), and b) a *Global DB(G)* having K global characteristics. The problem at hand is thus a *record matching* problem where products in local database are to be mapped to global characteristic values (e.g. 'category', or 'global brand' etc.).

Copyright © 2015 Tata Consultancy Services Ltd.

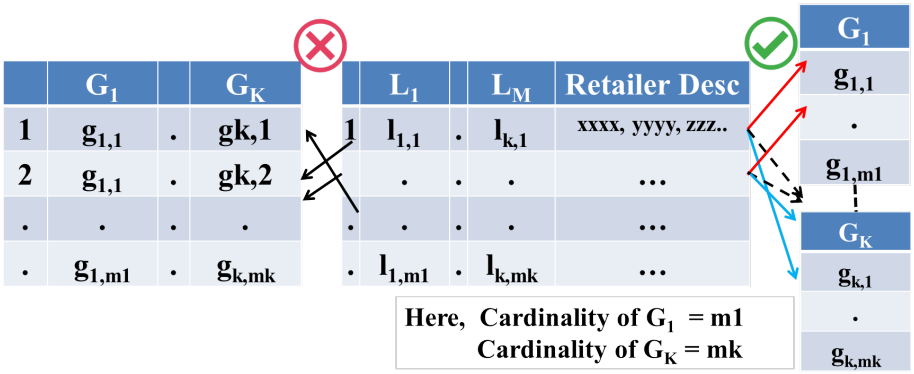

Figure 1: Local and Global Database

Note that our objective is only to reconcile performance metrics (such as volume sales and market-share) across databases for each global characteristic *independently*, e.g., sales vs market-share for each category, or alternatively each global brand, etc. We can achieve this by solving K different record matching problems, as shown in Figure 1: For each product, we shall predict each of the K global characteristics given local characteristics and retailer descriptions separately, as $\arg\max_j P(G_j|L_1, ..., L_M, D_i)$.

In this paper: a) We address the problem of automating attribute fusion across diverse data sources that do not share a common join key. b) We augment traditional, fundamentally unsupervised text-similarly techniques with supervised, Bayesian network models in a confidence-based ensemble for automating the mapping process. c) Our approach additionally delivers confidence bounds on its predictions, so that human annotation can be employed when needed. d) We test our approach in a real-life market research scenario. We also compare it with available techniques [3] and demonstrate that our approach outperforms FEBRL [3]. e) We illustrate how our approach has been integrated into a data-fusion platform[10] specifically designed to manage data-lakes containing disparate databases.

Related Work: Record linkage has been usually addressed via two categories of approaches, learning-based and non-learning based [6]. Learning-based approaches such as FEBRL[3] that uses SVM to learn a weighted combination of similarity matching techniques followed by unsupervised matching, MARLIN [1] uses similarity measures Edit Distance and Cosine and several learners. In non-learning based approaches, PPJoin+[11] is a single-attribute match approach using sophisticated filtering techniques for improved efficiency, and FellegiSunter[4] evaluates three of the similarity measures Winkler, Tokenset, Trigram. In[8], an ensemble approach of two non-learning algorithms Fellegi-Sunter and Jaro-Wrinkler has been presented for record-linkage. In contrast, we use a confidence based ensemble approach that combines supervised learning using a Bayesian network model together with a non-learning based textual model. Our approach also produces confidence bounds on the predictions that help to decide reliability of prediction.

2 Approach

Each product l in L has two kind of information (1) M Local characteristics and (2) Textual descriptions by retailers. In this section, we present our approach to predict the value of global characteristic G_j for each product in L. We use two different models for two different datasets (1) Supervised Bayesian Model (SBM) using local characteristics, and (2) Unsupervised Textual Similarity (UTS) using descriptions to compute probability of every possible state $g_{j,t}, t = 1, 2, ..., mj$ of G_j. Finally, we use an weighted ensemble based approach to combine the probabilities of both models to predict the value of G_j.

2.1 Supervised Bayesian Model

Approach to build SBM comprises of:(1) Network Structure Learning, (2) Parameter Learning, & (3) Bayesian Inference. For structure learning, we propose a novel technique of learning Tree based Bayesian Networks(TBN), whereas for parameter learning and Bayesian inference, we use the idea of [12] that performs probabilistic queries using SQL queries on the database of conditional probability tables.

TBN Structure Learning: Bayesian networks are associated with parameters known as conditional probability tables (CPT), where a CPT of a node indicates the probability that each value of a node can take given all combinations of values of its parent nodes. In CPTs, the number of bins grows exponentially as the number of parents increases leaving fewer data instances in each bin for estimating the parameters. Thus, sparser structures often provide better estimation of the underlying distribution [5]. Also, if the number of states of each node becomes high and the learned model is complex,

Copyright © 2015 Tata Consultancy Services Ltd.

Bayesian inferencing becomes conceptually and computationally intractable [7]. Hence, tree-based structures can be useful for density estimation from limited data and in the presence of higher number of states for facilitating faster inferencing. We employ a greedy search, and score based approach for learning TBN structure.

Given the global characteristic G_j and M local characteristics, we find set of top η most relevant local characteristics w.r.t. G_j using mutual information. We denote these η local characteristics by $Y^j(L)$. Further, we learn a *Tree based Bayesian Network(TBN)* on random variables $X = \{X_r : r = 1, 2, ..., \eta + 1\}$, where each $X_r \in X$ is either local characteristic $L_i \in Y^j(L)$ or global characteristic G_j

Chow et al. in [2] state that cross-entropy between the tree structures distributions and the actual underlying distribution is minimized when the structure is a maximum weight spanning tree(MST). So, in order to learn TBN structure, we first learn MST for the characteristics in the set X. We find the mutual information between each pair characteristics, denoted by $W(X_r, X_s)$. Further, we use the mutual information as the weight between each pair of characteristics and learn MST using Kruskal's algorithm.

$$TotalWeight(TW) = \sum_{r=1, Pa(X_r)\neq 0}^{\eta+1} W(X_r, Pa(X_r)) \tag{1}$$

By learning MST, order of search space of possible graphs is reduced to $2^{O(\eta)}$, from $2^{O((\eta)^2)}$. Using this MST we search for the directed graph with least cross-entropy, by flipping each edge directions sequentially to obtain 2^η directed graphs along with their corresponding TW calculated using Eq. 1. Graph with maximum TW (minimum cross-entropy) [7] is chosen as the best graphical structure representative of underlying distribution.

Parameter Learning and Inference: To learn the parameters of Bayesian Network(CPTs), for every product l in L we compute the probabilities $p^l_{j,1}, p^l_{j,2}, ..., p^l_{j,m_j}$, for every state of G_j, given the observed values of local characteristics in the Bayesian network, using an approach described in[12]. Here, CPTs are learned from the data stored in RDBMS and all queries are also answered using SQL.

2.2 Unsupervised Text Similarity

In this section, we present UTS approach to compute the probability $q^l_{j,t}$ of each possible state of the global characteristic G_j using retailer descriptions. Consider each product l in L has r_l descriptions and for each description $d_{l,r}$, where $r = 1, 2, ..., r_l$, we find n-grams of adjacent words. Let $N_l = \{n^l_v, v = 1, 2, ...\}$ be the set of n-grams of all descriptions, where f^l_v be the frequency of each n^l_v defined as a ratio of the number of descriptions in which n^l_v exists to the r_l.

For every state $g_{j,t}$ of G_j, we find the best matching n-gram from the set N_l by calculating Jaro-Wrinkler distance between $g_{j,t}$ and every $n^l_v \in N_l$ and choose the n-gram, say $n^l_{v,t}$, with the maximum score $s^l_{j,t}$. Further, multiply the scores $s^l_{j,t}$ with the frequency of $n^l_{v,t}$ to get the new score i.e., $S^l_{j,t} = s^l_{j,t} \times f^s_{l,t}$. Finally, we convert each score $S^l_{j,t}$ into the probability $q^l_{j,t}$ by using softmax scaling function.

2.3 Ensemble of models

In ensemble approach, we first find confidence of each prediction in both the cases(SBM and UTS) and then use these confidence values as weights for weighted ensemble. Given the probability distribution $\{p^l_{j,t} : t = 1, 2, ..., m_j\}$ for the values of G_j using SBM model, we find the confidence corresponds to each probability as

$$C(p^l_{j,t}) = 1 - \sqrt{\sum_{t'=1}^{m_j} (p^l_{j,t'} - h^l_{t'}(t))^2}, t = 1, 2..., m_j \tag{2}$$

where $h^l_{t'}(t)$ is the ideal distribution, which is 1 when $t = t'$ and 0 otherwise. Similarly, we can find the confidence $C(q^l_{j,t})$ of each probability $q^l_{j,t}$.

With the given probability dist. and the confidence values from both models, we take weighted linear sum of two probabilities to get the new probability distribution over the states of G_j: $P^l_{j,t} = C(p^l_{j,t}) \times p^l_{j,t} + C(q^l_{j,t}) \times q^l_{j,t}$, $t = 1, 2, ..., m_j$ and we choose the value of G_j for maximum $P^l_{j,t}$.

CoP: For every prediction, we assign the confidence value called confidence of prediction (CoP). CoP is a measure that helps to decide whether the predicted value is trustworthy or not. Given the probability distribution $\{P^l_{j,t} : l = 1, 2, ..., m_j\}$ for the values of g_j, we calculate the CoP of the predicted value $g^l_{j,t}$ of G_j by using Equation 2.

3 Experiments and results

We present the accuracy of our predictions on a real-life dataset from a global market research organization. We set a threshold τ on CoP, and predictions with a CoP $< \tau$ are routed for human annotation. We also measure the accuracies of our predictions for different values of τ.

Copyright © 2015 Tata Consultancy Services Ltd.

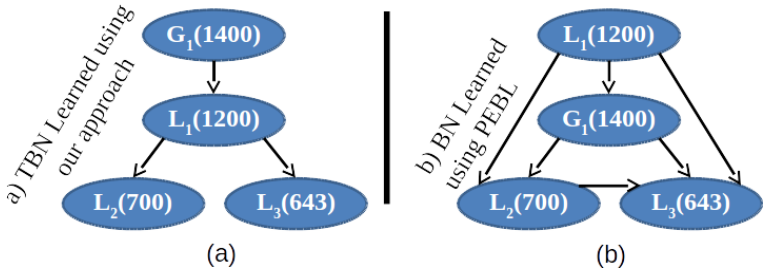

Figure 2: Structure learned of G_1 using two different approaches

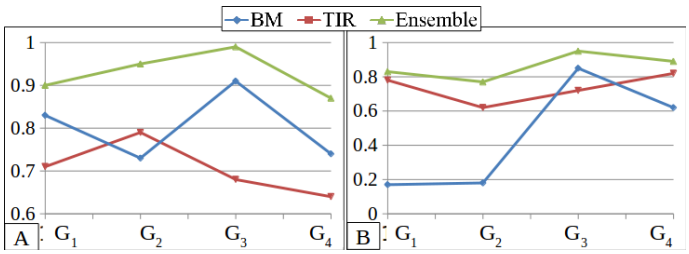

Figure 3: x-axis: global characteristics, y-axis: A) Predictive accuracy for Case-1, B) Predictive accuracy for Case-2

Data description: We have data for carbonated drinks of 26K unique products from a single geography, contained in two datasets: a) **Local DB**: It contains 26K products with each product having 49 local characteristics, where cardinality of local characteristics varies from tens to thousands. It also contains descriptions of products given by retailers of that product, where number of descriptions of a single product varies from tens to hundreds. b) **Global DB**: It contains four global characteristics with cardinality varying from tens to thousands.

Data Preparation: We predict four global characteristics G_1, G_2, G_3, and G_4 for two cases, with varying ratio of split between training, validation and test datasets. ***Case-1*** (60:20:20) has 60% training, 20% validation and 20% test and ***Case-2*** has this ratio as 20:20:60. **NOTE:** While Case-1 uses a traditional split of training vs testing data, Case-2 is more realistic, since in practice preparing a training data by manual data labeling is costly: For example, we would like to 'onboard' a data from a particular dataset by manually annotating only a small fraction (e.g. 20%) of records and automate the remainder or we might like to board data from one organization (e.g. retailer or distributor) in a particular geography in the hope that data from remaining sources in that geography share similar local characteristics, eliminating manual annotation for a large volume of data. To simulate this practical scenario, we used the first few records from the local dataset, which happened to contain only 10% or so of the total possible values of each global attribute.

For SBM, η relevant local characteristics was chosen for every G_j. Figure 2, compares the Bayesian network structure learned using our approach and another learned using an open source python library Pebl [9], for the global characteristic G_1. Clearly, network obtained using Pebl (Figure 2(b)) is more complex as compared to ours 2(a), as the size of CPTs of these are of the order of a) 1200×1400 and b) $1200 \times 1400 \times 700 \times 643$ respectively.

Figure 3, shows the prediction accuracy of four global characteristic for Case-1 and Case-2 respectively. Here, the accuracy is a ratio of correctly predicted products to the total number of products. In Case-1, accuracy of Ensemble model is in the range of 85 to 99% and it outperforms both SBM and UTS for all four global characteristics.

Baseline Comparison: We also compared our approach with record matching method implemented in a framework called FEBRL[3]. For attribute matching, we tried three similarity measures winkler, tokenset, trigam and show the results with winkler which outperforms the rest. We tested this approach for the Case-1 on the smaller dataset (5K products). Table 1, shows the comparison of the prediction accuracy of four global attributes using our Ensemble approach and FEBRL. This suggests that our approach outperforms and also shows that accuracy of FEBRL decreases for high cardinality global attributes. FEBRL did not work on, 26k products, on a machine with 16GB RAM, Intel Core i7-3520M CPU 2.90GHz* 4, 64 bit. We did not try the blocking method as main motive of our problem is to improve accuracy of prediction, and not the time complexity. **Case-2** (Figure 3-B), naturally renders the SBM less accurate, since the training data contains only 10% of possible states of each global characteristic. However, it is compensated by the performance of UTS, which searches the target set of global attribute values from the retailer descriptions. Combining these models using our Ensemble model the accuracy of four global characteristics reaches 78 to 93%.

CoP Threshold for human annotation: We define three categories: a) **P-C:** Number of products *predicted correctly* by our approach for which CoP $> \tau$. b) **P-I:** Number of products *predicted incorrectly*, for which CoP $> \tau$. c) **NP:** Products

Copyright © 2015 Tata Consultancy Services Ltd.

Global att.	Num of states	FEBRL(winkler)	Ensemble
G_1	107	86%	93%
G_2	154	57%	95.2%
G_3	3	99.3%	99.2%
G_4	13	95.4%	99.4%

Table 1: Comparison of our approach with FEBRL

Global	Case-1				Case-2			
	τ	P-C	P-I	NP	τ	P-C	P-I	NP
G_1	0.5	92%	4%	4%	0.6	82%	7%	11%
G_2	0.6	81%	7%	12%	0.65	74%	10%	16%
G_3	0.7	96%	1%	3%	0.7	96%	1%	3%
G_4	0.8	86%	3%	11%	0.8	85%	4%	11%

Table 2: % of products in each category on Validation set

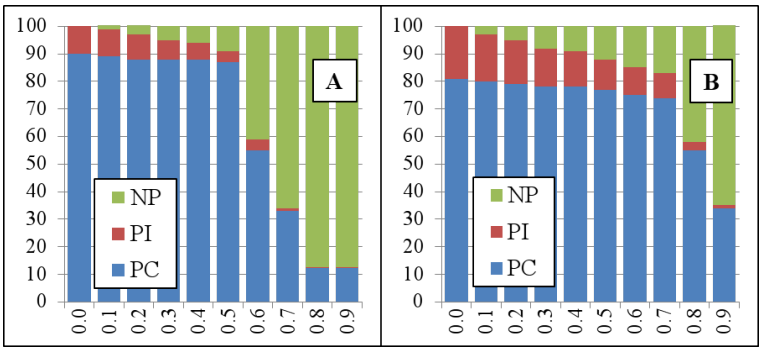

Figure 4: % of Products in each category for different values of τ on test data for G_1 in A) Case-1 and B) Case-2

which we choose *not to predict*, i.e., products with CoP $\leq \tau$. We select τ in order to maximize P-C and minimize P-I category, while not increasing NP so much that exercise becomes almost entirely manual. Since products in the P-I category are more costly for a company as compared to NP category, we give more weight to P-I while learning τ. Table 2, shows the percentage of products in each category (P-C, P-I, NP) on validation set along with the threshold τ values for both cases. It shows that for given τ, percentage of products in P-C category is in the range of 81-96% for Case-1, whereas, it ranges from 70 to 96% for Case-2. Also, the average percentage of products in P-I category is only around 5%. These numbers establish that CoP is a good measure for reliability of predictions. Figure 4, shows the variation in the percentage of products in test set of each category with respect to threshold value τ for both Case-1 and Case-2, for the global characteristic G_1. It validates the optimal values of τ learned using validation set, 0.5 for Case-1 and 0.6 for Case-2.

The process of aggregate analysis, comparing global market share and sales of product categories is carried out in our platform ***iFuse***[10] (Figure 5). Figure 5(a),(b) shows the data tile and cart view of iFuse representing the attributes of the local DB and global DB to be linked together. Figure 5(c) shows the tile view of the attributes obtained after mapping of local DB to global attribute, here GLO BRAND via ensemble approach, thereby enabling the **join** of local sales and global market share via common global attribute, GLO BRAND (Figure 5(d)). Figure 5(e) shows aggregate analysis of different products via motionchart.

4 Conclusion

We have addressed a particular class of record-linkage problems where disparate databases need to be fused in the absence of matching keys for the limited purpose of aggregate analysis. Our ensemble approach combines supervised Bayesian models with unsupervised textual similarity, and also returns confidence along with each prediction. We submit that our approach is likely to be applicable for similar instances of record-linkage in a wide variety of applications, even while attempting to fuse data from external sources, such as social media, sensor data etc.. Such scenarios are becoming increasingly common as the *data lake* paradigm is gradually replacing the traditional data-warehouse model, driven by the availability and accessibility of external 'big data' sources.

Copyright © 2015 Tata Consultancy Services Ltd.

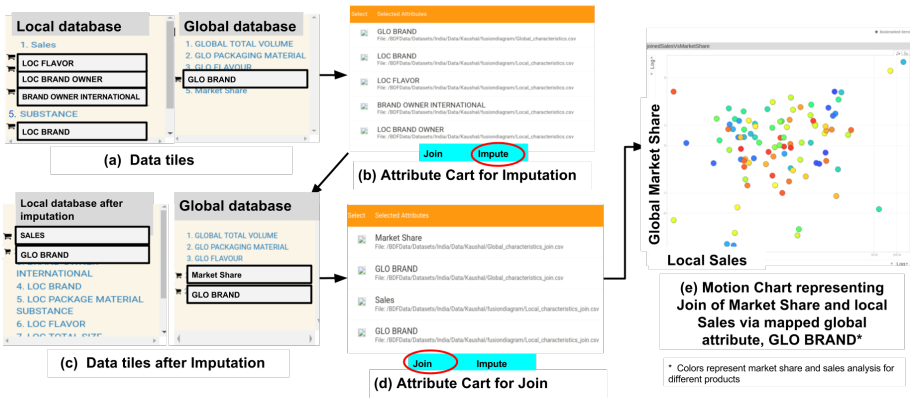

Figure 5: Figure showing aggregate analysis of global market share and local sales done using our platform.

References

[1] Bilenko, M., Mooney, R.J.: Adaptive duplicate detection using learnable string similarity measures. In: In ACM SIGKDD int. conference on Knowledge discovery and data mining. ACM (2003)

[2] Chow, C., Liu, C.: Approximating discrete probability distributions with dependence trees. Information Theory, IEEE Transactions on (1968)

[3] Christen, P.: Febrl: a freely available record linkage system with a graphical user interface. In: In 2nd Australasian workshop on Health data and knowledge management-Volume 80. Australian Computer Society, Inc. (2008)

[4] Fellegi, I.P., Sunter, A.B.: A theory for record linkage. Journal of the American Statistical Association (1969)

[5] Koller, D., Friedman, N.: Probabilistic graphical models: principles and techniques. MIT press (2009)

[6] Köpcke, H., Thor, A., Rahm, E.: Evaluation of entity resolution approaches on real-world match problems. Proceedings of the VLDB Endowment (2010)

[7] Lam, W., Bacchus, F.: Learning bayesian belief networks: An approach based on the mdl principle. Computational intelligence (1994)

[8] Poon, S., Poon, J., Lam, M., et al.: An ensemble approach for record matching in data linkage. Studies in health technology and informatics (2016)

[9] Shah, A., Woolf, P.: Python environment for bayesian learning: Inferring the structure of bayesian networks from knowledge and data. Journal of machine learning research: JMLR (2009)

[10] Singh, K., Paneri, et al.: Visual bayesian fusion to navigate a data lake. In: Information Fusion (FUSION), 2016 19th Int. Conference on. ISIF (2016)

[11] Xiao, C., Wang, W., Lin, X., Yu, J.X., Wang, G.: Efficient similarity joins for near-duplicate detection. ACM Transactions on Database Systems

[12] Yadav, S., Shroff, G., Hassan, E., Agarwal, P.: Business data fusion. In: Information Fusion (Fusion), 2015 18th International Conference on. IEEE (2015)

Copyright © 2015 Tata Consultancy Services Ltd.